

# Spatio-Temporal Small Worlds for Decentralized Information Retrieval in Social Networking

Georg Groh  
TU München  
Faculty for Informatics  
grohg@in.tum.de

Florian Straub  
ETH Zürich  
Inst. of Cartography and  
Geoinformation  
straubf@ethz.ch

Benjamin Koster  
TU München  
Faculty for Informatics  
koster@in.tum.de

## ABSTRACT

We discuss foundations and options for alternative, agent-based information retrieval (IR) approaches in Social Networking, especially Decentralized and Mobile Social Networking scenarios. In addition to usual semantic contexts, these approaches make use of long-term social and spatio-temporal contexts in order to satisfy conscious as well as unconscious information needs according to Human IR heuristics. Using a large Twitter dataset, we investigate these approaches and especially investigate the question in how far spatio-temporal contexts can act as a conceptual bracket implicating social and semantic cohesion, giving rise to the concept of Spatio-Temporal Small Worlds.

## Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

## Keywords

Collaborative (Geographic) Information Retrieval, Spatial Context, (Geo) Social Networks, Spatial Context, Human Search, Small World Networks, Data Analysis, Information Needs.

## 1. INTRODUCTION

Social Networking (SN) and Decentralized Social Networking (DSN) [53] as a future variant of SN is extensively used to build rich personal and interpersonal information spaces. Furthermore, the increased access of SN-platforms via mobile devices such as smartphones (giving rise to new paradigms such as (context-aware) Mobile Social Networking (MSN)) introduces a steeply growing permeation of these information spaces with explicit spatial context. Thus, besides social contexts such as ‘friendship’ relations, spatio-temporal contexts and their interrelations with social contexts are also available and extensively used in modern (M)SN platforms.

These upcoming SN paradigms allow users more and more to employ special forms of information retrieval, akin to traditional human information seeking behavior based on the real social network of society (‘Human IR’) which, besides semantic context, also uses social and spatio-temporal context (see also [52]).

Inspired by this behavior, the question now arises how alternative IR services for SN may be constructed that effectively make use of social, semantic, and spatio-temporal contexts and their interrelations.

Pursuing this research question, the remainder of this paper is structured as follows. After a brief discussion of the

relation between context and information needs, we address Human IR and wayfinding in social networks. We then introduce the concept of Spatio-Temporal Small Worlds for IR in Social Networking as well as a respective architecture based on personal information agents. The following main part of the paper empirically investigates the concept of Spatio-Temporal Small Worlds and the suitability of the principles guiding alternative IR processes inspired by Human IR, using social search, semantic search and spatio-temporal search and here especially the suitability of spatio-temporal embedding as a contextual bracket using a large Twitter dataset.

This paper is an extended version of the content of the paper [13]. Elements of this text also appear in the thesis [14].

## 2. RELATED WORK AND FUNDAMENTAL CONSIDERATIONS

### 2.1 Context and Unconscious Information Needs

In [40] adequate characterizations of *relevance in information retrieval* (IR) and especially *qualifications of information needs* that a user of IR has in view of a ‘problematic situation’ [7], [8], [40] are investigated. In this regard, the concepts query, request, perceived information need (PIN), and real information need (RIN) are considered as central. The query is a formalization of a request which, in turn, is a natural language expression of a PIN. The PIN is the information need that a user subjectively *perceives* in the problematic situation. The RIN may e.g. be defined via the entirety of information that is ‘objectively’ relevant for the solution of the problem, thus *extensionally defining* the ‘problem’ in ‘problematic situation’ through the RIN. ‘Objectively’ may e.g. be determined by the intersection or union of the assessed RIN by the fictional set of all human experts for the problem.

During the IR process the user then consumes or partly consumes the results, uses his assessment of relevance judgments, corrects his PIN, formulates a new query and so on, giving rise to a *circular IR process* (see e.g. [6]). A user will explore the space of information relevant to the RIN by repeated executions of the aforementioned IR cycle, iteratively re-shaping his PIN, and enlarging the set of acquired information.

Our notion of *conscious information need* corresponds to *perceived information need* (PIN) in [40] and our notion of *unconscious information need* encompasses the *real information*

tion need (RIN) in [40]. In IR, the term unconscious information need is justified because the user is not consciously aware of information needs in  $RIN \setminus PIN$  (that are in RIN but not in PIN) in a ‘problematic situation’. However, our notion of *unconscious information need* also encompasses *an unspecific readiness to accept ‘interesting’ information*. Unless artificially defining some ‘background problematic situations’, *ongoing readiness to accept welcomed information that does not correspond to a ‘problematic situation’* (and thus not to a RIN or PIN) is not represented in the schema of IR relevance. This case is simply not covered by the concept of *information retrieval*, where a problematic situation induces a concrete information need which in turn finally induces a query. Examples for such a form of *unconscious information need* correspond to e.g. a user reading ‘something interesting’ on a news-feed or is being told ‘something interesting’ by a friend, etc. Thus information may be delivered to a user that the user has no a priori *perceived* information need for, and which the user has not explicitly asked for via a query or filter, but that he / she nevertheless judges as ‘interesting’. This kind of information is usually pro-actively delivered by awareness services, or news services, or by direct communication services [14].

*Context* and especially *social context* may be used to provide a *relevance bracket* for this ‘interesting information’ that is delivered to a user by such services by e.g. narrowing the visualizations of social network dynamics to the network neighborhood or spatio-temporal neighborhood of a user [17], using social filtering to deliver horizon broadening recommendations [16], or using social contexts to specify suitable audiences for certain information [15].

The *contextual relevance bracket* is a means to anticipate or *induce relevance via context* in these proactive services [14]. *Incorporation of context*, especially of social and spatio-temporal context, can be *especially useful for information retrieval* e.g. by aiding the user in exploring the space of relevant information items / in expanding the PIN, especially in relation to problems for which the RIN is hard to determine. This aid can be achieved by seeding the IR cycle with new motives especially beyond the PIN while providing a certain contextual bracket for relevance (in contrast to e.g. randomly choosing the seeds) as Figure 1 illustrates. In con-

petence gains, or perspectives appearing after new IR cycles and thus  $PIN \approx RIN$  after ‘sufficiently’ many IR cycles, the ‘problematic situations’ for which the RIN is hard to determine might not exhibit this saturation effect, either because the problem’s definition is not precise enough or because the space of information items relevant to the RIN is very large.

Traditional *Context-Sensitive Information Retrieval* is usually focused on using types of context such as query histories or implicit feedback on the results to a query (e.g. via click-analysis or eye-tracking) to improve relevance of the immediately retrieved results in view of a given query (see e.g. [47]). However it is usually limited to the PIN expressed in the query, because more general contextual brackets (e.g. induced by social context) that would be able to deliver the contextual seeds mentioned before are missing or not regarded. E.g. including seeds from the information spaces of other competent people determined via (besides the query) also taking social context into consideration, may improve the exploration of the RIN, especially in those cases where the boundaries of the RIN are hard to determine precisely.

## 2.2 Human IR and Wayfinding in Social Networks

If *long-term social contexts* in the form of *social networks* are used to provide *contextual brackets* for information retrieval services in SN / MSN, it is important to review the basic results of *decentralized routing and searching* in these networks [33].

In 1967, *Milgram’s experiment* [39] showed that decentralized routing in social networks is possible and that the path lengths involved were small. Watts and Strogatz [50] were able to provide a network model for such *Small World networks*, which did not only explain their short mean average path length but also their high clustering coefficient (the network theoretic measure for triadic closure), a crucial property of social networks. The *Watts-Strogatz model* is based on a toroidally, regularly linked graph, where edges are randomly redirected with a certain probability (shortcuts). These constructive elements generate the local cluster structure and short mean average path length. [50].

While such models were able to explain the basic *structure of social networks*, the actual explanation of the Milgram experiment, the question of how *decentralized wayfinding or routing can actually be accomplished*, was investigated by Kleinberg [29]. In his variant of the Small World model, starting from a regularly linked network on a grid, the random distant re-connections of a node  $a$  to a node  $b$  were established with a probability  $d(a, b)^{-\alpha}$ . He was able to show that for  $\alpha$  corresponding to the dimension of the grid, a decentralized (local knowledge only) routing algorithm, always choosing the node located closest to the target node as the next node, is sufficient to produce ‘sufficiently’ short expected delivery times, polynomial in  $O(\log(n))$ , where  $n$  is the number of nodes in the network. Refinements of this model in view of more realistic *geographic distributions of friendship relations* on the earth’s surface were investigated by [34], suggesting a different geographic connection probability distribution and empirically finding a different value for  $\alpha$ , but confirming that the simple greedy local routing algorithm still leads to efficient delivery. This confirms that for *efficient decentralized geographic routing in social networks*, the nodes (actors) of the network need to be *spatially embedded* (e.g. have a known center of life) and each forwarding

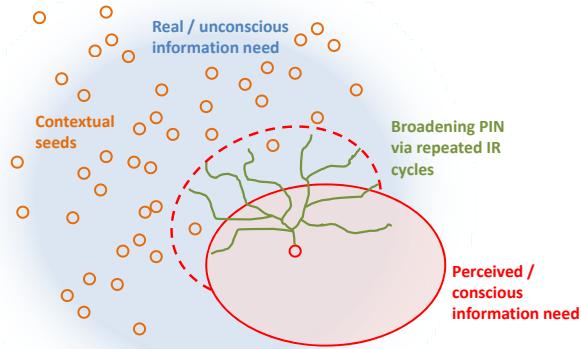


Figure 1: Defining unconscious information need [14]

trust to well defined problems, which may exhibit a natural saturation effect in view of new information, insights, com-

actor needs to have a *cognitive model of this spatio-temporal context*.

More generally, besides spatial proximity other types of contextual metrics such as other long-term social contexts (e.g. occupation or hobbies) may as well be chosen to select the next node. The greedy local social search will select as the next node the node closest to the target node according to the given metrics (see e.g. [33]).

Parallels exist between using general context information for decentralized routing and the way *social information retrieval* is accomplished in *human societies*, which in turn has obvious *commonalities with SN / MSN*. In '*Human IR*', a question formalizing a PIN is 'routed' to persons which presumably dispose of the required information in their (not necessarily properly explicated) information spaces. The resulting routes need to be 'socially resilient' enough (e.g. in the sense of Granovetter's strong ties [12]) to support the actors en-route agreeing to process the query and to support routing the retrieved information back to the questioner. At the same time the routes must contain enough weak ties (in Granovetter's sense) to convey new information or provide access to otherwise hardly reachable parts of the network via weak tie shortcuts in the sense of [50] [11].

As reviewed in [52], *human information seeking behavior* often use context e.g. social context to determine actors who could be asked, especially if the problem situation and the PIN is poorly defined ([52]). Actors facing an informational problem will, besides the PIN ( $\hat{=}$  WHAT), evaluate *all types of contexts*, their interrelations and their relations to the PIN, in order to render their PIN more precise, expand their PIN towards the RIN and ultimately collect enough information to solve their problem (see [18] for a more elaborate discussion). For the discussion, types of contexts will be represented by other interrogative pronouns such as WHO (pointing to social context), WHERE and WHEN (pointing to spatio-temporal context). *Vice versa*, the asked persons may also use contextual knowledge to select appropriate information for the questioner, which may also include information that is not strictly relevant to the query but relevant to the PIN or even RIN of the questioner. Thus relevance may also be induced by the asked actor *via contextual knowledge*. As an example consider the question "How do I search for certain terms while I browse a text-document with UNIX 'more' ?". As an *expert*, a person might answer "Use the '/' character and enter the term". As an *expert* and friend the answer may include "Besides: use 'less' instead of 'more'! It has a number of advantages". As an *expert and close friend* the answer may include "Besides: I give You the advice to quit using UNIX! A Mac will suit Your needs and the needs of Your wife much better. It provides more comfortable means to view and search text-files while still retaining 'less' and 'more' if desired", using social context and the questioner's individual context.

In terms of *long-term social context*, *Human IR* 'uses' the main *characteristics of small world networks* to search in the complex network of distributed information spaces and context-elements for the right information: actors are able to reach experts (and their information spaces) via short expected path lengths and the highly clustered structure ensures that each actor has a large number of routing options. Suitable interdependent contextual metrics (Semantic (WHAT), social (WHO) or spatio-temporal (WHERE + WHEN)) allow efficiently *navigating the space*.

### 3. SPATIO-TEMPORAL SMALL WORLDS FOR IR IN SOCIAL NETWORKING

The question now arises, how we can employ these considerations and the considerations of the preceding section to *construct an alternative information retrieval service for SN / MSN*. While the complex socio-psychological mechanics of amalgamating and evaluating the interdependencies of WHO  $\leftrightarrow$  WHERE  $\leftrightarrow$  WHAT  $\leftrightarrow$  WHEN in Human IR in view of searching the distributed information spaces in a context sensitive way are too intricate to model directly, *spatio-temporal embedding may act as a reference point* and a means to naturally encode these interdependencies between the various forms of context for a respective IT model.

#### 3.1 Spatio-Temporal Small Worlds

A *social spatio-temporal small world* may be defined as a social network, where the actor-nodes are *spatio-temporally embedded* into space-time via their current center of life (compare previous section). The relations correspond to directed long-term social relations of various types. We have seen that spatial distance metrics (and via using a current time-frame thus also spatio-temporal distance metrics) allow efficient decentralized routing. We assume that spatio-temporal distance metrics can thus also serve as one key means for a successful search for information in the social spatio-temporal small world part of the complex network of distributed information spaces and context-elements described in the previous section. 'Successful' implies that the information found is relevant in view of a user's RIN especially in those cases where the RIN is hard to demarcate (see discussion in subsection 2.1). Another argument for using spatio-temporal distance metrics as a means to naturally encode interdependencies between the various forms of context or other metrics is that the studies of Kleinberg [30] and Liben-Nowell [34] imply that in a social spatio-temporal small world, *spatio-temporal closeness is probabilistically correlated with social closeness*.

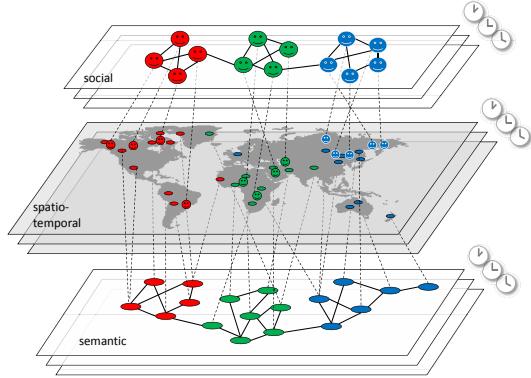
A *semantic spatio-temporal small world* may be defined as a network of information items (e.g. documents) that are *spatio-temporally embedded* into space-time via certain criteria. Either the information item's meta-data contains an *explicit spatio-temporal embedding* or *implicit spatio-temporal embedding*, e.g. explicated spatially via geo-parsing (see e.g. [32] [28]) and geo-coding (see e.g. [28]) of the found named entities (see e.g. [41]). A third case applies if the information item is spatio-temporally embedded in the *same spatio-temporal location(s)* as the actor whose information space this item is associated with.

The *first mode of edges* of this network are the links indicating semantic relatedness of the items (e.g. HTTP links). The corresponding network has small world properties [26]. The *second mode of edges* relates items, whose 'owners' are linked in the social spatio-temporal small world, which also gives rise to a network with small world properties.

As previously discussed, *social closeness is probabilistically correlated with spatial (and implicitly spatio-temporal) closeness* [33] [46].

Studies by Brent Hecht [22], [23], [21], [24] and others (e.g. [35]) imply that in a semantic spatio-temporal small world, *spatio-temporal closeness is probabilistically correlated with semantic closeness* to a certain extend, which is also expressed as a statistical tendency in (so-called) Tobler's first law of Geography: "everything is related to everything else,

but near things are more related than distant things" [49].



**Figure 2: Spatio-temporal embedding of small worlds: how spatio-temporal embedding maintains social and semantic closeness properties as a statistical tendency [14][18].**

Figure 2 visualizes social and semantic spatio-temporal small worlds and illustrates the maintenance of social and semantic closeness via spatio-temporal embedding.

*Social closeness is also probabilistically correlated with semantic closeness.* Homophily (the tendency of similar people to associate with each other, contributing to triadic closure) [38] and Peer Influence (the influence of persons which are directly linked in the social network) [10] can be prominently attributed for the local homogeneity in terms of information spaces of social groupings. The correlation between social closeness and semantic / topical closeness is also supported by other studies such as [5] and indirectly by [16].

Thus in view of decentralized search of relevant information in the complex network of distributed information spaces and context-elements which is characteristic of SN / MSN, we assume that social spatio-temporal small worlds and semantic spatio-temporal small worlds may act as a simple model of this complex network of context elements and spatio-temporal metrics may aid the decentralized search because of implicitly representing interrelations between spatio-temporal, social, and semantic relatedness.

Based on these considerations and models and the principles of Human IR, the *study* [18], proposed a new *context-aware, agent-based, federated approach to information retrieval in decentralized SN / MSN* in order to investigate limits and chances of using spatio-temporal embedding and its implicit 'conservation' of semantic and social context as a contextual bracket. Besides the discussion of the last sections, the design decisions in this study were supported by a number of observations such as the ever growing availability of context in SN and especially MSN, the importance of the paradigm of Distributed Social Networking [53], the problems that the Hidden Web especially in connection with access protected SN / MSN information spaces generates for traditional search engines [20], or the obvious parallels that searching in SN / MSN has to Human IR.

### 3.2 An Architecture based on Personal Information Agents

The *architecture* of [18] is based on *personal information agents* associated with *spatio-temporally embedded social ac-*

*tors* (users, companies, SN-platforms etc.), which contextually decide upon the execution of another actor's query in relation to the asked actor's information space. The agents are able to answer these queries in a context sensitive way, using techniques from Context-Sensitive IR and their expertise on their own information spaces. Each actor maintains socio-semantic links to other spatio-temporally embedded actors in form of topic specific *expert-links*, thus implementing a special form of social spatio-temporal small world. Furthermore, each actor publishes a selection of his / her expert-links and a set of spatio-temporally embedded *expertises*, summarizing content fields from the actor's information space (thus contributing to a special form of a semantic spatio-temporal small world). The spatio-temporal embedding of expert-links and expertises ('*knowledge flags*') follows the three step process discussed above. These knowledge flags are published in a decentralized *spatio-temporal Peer-to-Peer index*. If an actor issues a query which cannot be answered from his own information space, a social search is performed using the actor's expert links. If this search also fails, the spatial index is queried using the spatio-temporal embedding of the query, with the embedding following the three step process: e.g. if the query does not contain a spatio-temporal reference, the spatio-temporal reference of the questioner (see subsection 2.2) is used. The search delivers a number of knowledge flags which the questioner's agent then further evaluates by asking the related other agents. The system thus combines elements of social search (via expert-links), semantic search (local IR-systems) and spatio-temporal search (implying social and semantic contexts to a certain extent as explained above).

Compared to e.g. Peer-to-Peer (P2P) IR systems, were e.g. an index is distributed over the passively protocol-executing peers in a P2P network (see e.g. [48] for a hybrid document- / index-distribution approach), and thus in most cases basically 'merely' distributes a conventional IR system over a P2P network, this architecture uses the actor's agent's local IR systems to locally decide upon relevance. The agents are thus able to take into account the (e.g. social) context of the query and the querying agent / its user, thus being able to optimize contextual relevance and decide upon access [15] to control information flows, ensure privacy or even employ information markets [15]). Furthermore, they are able to pro-actively keep their published knowledge flags up-to-date.

The small world structure of the networks involved ensures that the expert-links, the comparatively coarse semantic mapping of the agent's information spaces in form of the expertises, and with the comparatively coarse implicit conserving of semantic and social contexts via spatio-temporal embedding is sufficient to deliver enough contextual seeds to reach enough competent agents which can then either employ their local IR systems to deliver contextually relevant items or use the private parts of their expert link list to further forward the query if the questioner's context is matching (e.g. if the corresponding user is a friend) resembling Human IR.

## 4. STUDY

### 4.1 Methodology

Some elements of the architecture (such as the specially designed spatio-temporal P2P Quad-Tree) were evaluated

using a dataset based on spatially referenced Wikipedia articles, demonstrating their practicability (see [19],[31]). Despite not disposing of a full implementation and evaluation scenario involving the necessary large number of actors and sub-systems, another evaluation step that can be taken is to evaluate the suitability of the principles guiding the architecture's IR process inspired by Human IR, using social search, semantic search and spatio-temporal search and here especially the suitability of spatio-temporal embedding as a contextual bracket for this type of IR, implying social and semantic contexts to a certain extent as explained above. For this evaluation, a data-set is required that contains real association of users and information items as well as realistic locations of users and explicit spatio-temporal references of their information items, as well as a social network exhibiting characteristics of the expert-link network proposed in the architecture. The micro-blogging service Twitter [4] with his network of followers, significant share of mobile usage and thus a large share of explicit spatial embeddings, and the free availability of the data is a suitable evaluation ground. We will now discuss some results of this evaluation.

## 4.2 Dataset

A dataset from Twitter was downloaded in June and July 2010, using the Twitter API [4]. The Tweets and Re-Tweets which were non-English (which was decided using the approach described in [9], employing an ML classifier using language specific n-gram statistics) were discarded. The remaining (Re-)Tweets were Porter-stemmed [44] and stop-words were removed. Of the Re-Tweets, only the additional content without 're-citing' the original Tweet was regarded. An undirected social network between the users was induced by establishing an edge if at least one @Reply or @Mention [4] (roughly corresponding to a direct message) was exchanged between the respective users. Of this social network, the largest connected component was chosen, and the rest of the users and their Tweets and Re-Tweets discarded. We downloaded the complete information from 43973129 Tweets and Re-Tweets, of which 9725514 were explicitly geo-coded. 3323803 of these geo-coded entities were associated with the largest connected component of our social network and finally considered. Of the 6887632 users in the dataset, 670271 were explicitly geo-coded and 160690 of these belonged to the largest component of the social network that we considered.

Users were spatially embedded via the geo-location of their last available explicitly geo-located (Re-)Tweet. (Re-)Tweets not explicitly spatially embedded (via geo-coordinates) were embedded with a simple geo-parsing approach, analyzing the strings denoting the location and subsequently using the MetaCarta geo-coding service [3]. If this process failed, the geo-location of the Wikipedia articles corresponding to the tags of the respective (Re-)Tweet, were used, using the Wikapedia API (see previous section). If that fails, the location of the authoring user was used. Locations were subjected to very small (uniform distribution in [-0.1,0.1] decimal degrees) random deviations to avoid mapping many entities to the exact same location which would result in overcrowding peers with respect to the Quad-Tree based spatio-temporal index which was used in the evaluation environment for the experiments.

## 4.3 Interrelations between Spatio Temporal, Social and Semantic Contexts

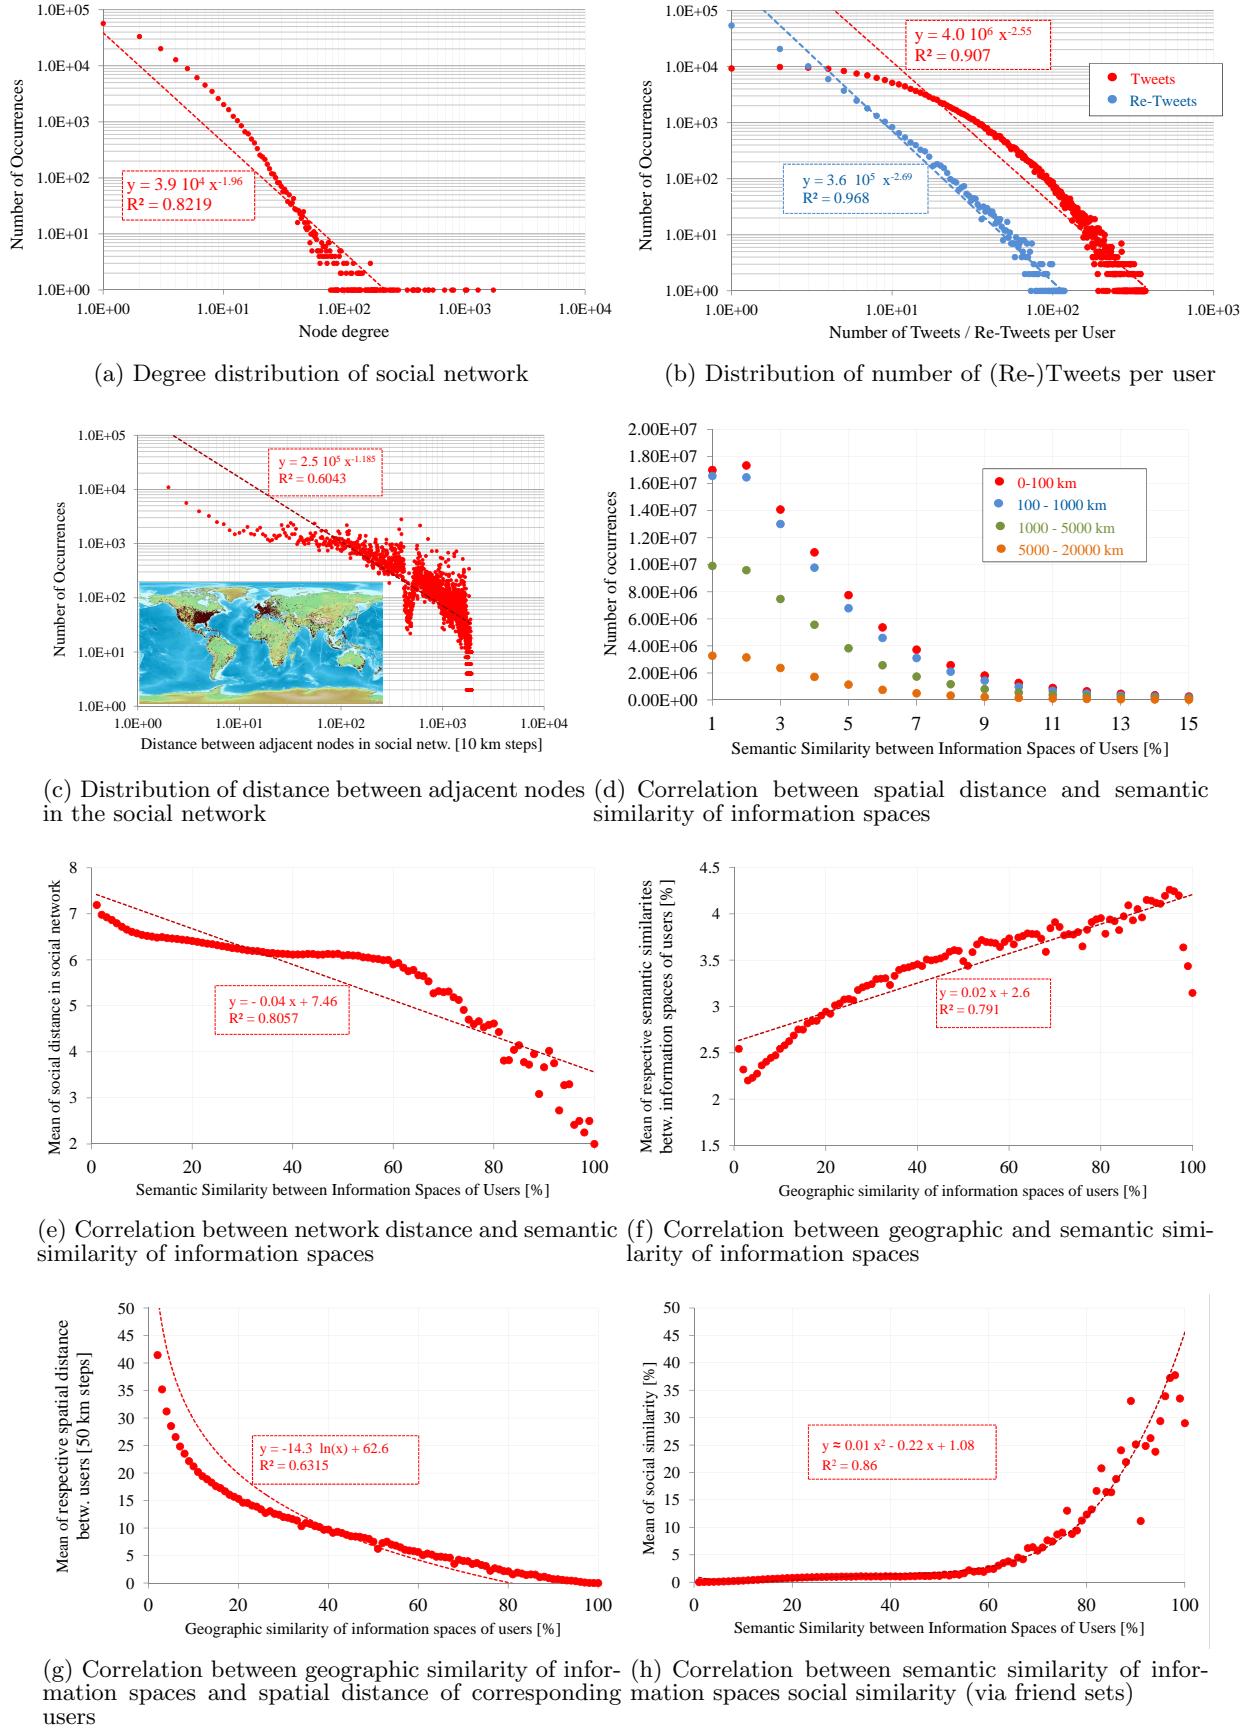
The *social network*'s mean average path length was 6.92 (a random graph with the same number of nodes, which was computed with the help of the JUNG framework [43] yielded a value of 8.96), and the average clustering coefficient [50] of the social network was 0.046 (corresponding random graph: 0.000014). We see that although the average clustering coefficient on SN platforms is usually higher by a factor of  $> 4$  (e.g. [51] report an average 0.164 for their early 2009 crawl of several sub-networks of Facebook with an overall number of nodes of  $\approx 10^6$ ). The numbers indicate that the present network can still be regarded as having small world properties.

*Figure 3* shows statistical properties of the dataset and correlation effects that support the mutual implication of social, semantic and spatial closeness which represents a basis for the proposed IR architecture. *Sub-Figure 3(a)* shows the degree distribution of social network which roughly follows a power law. This fact and the deviations from the exact power law distribution coincide with the findings in [42] [34]. Together with the previously discussed values for the mean average path length and clustering coefficient shows that the social network of actors in the data-set can indeed be assumed to be a realistic small world social network.

*Sub-Figure 3(b)* shows a distribution of the number of Tweets and Re-Tweets per user which, in our experiment simulate the information spaces of the users. While the Re-Tweet distribution follows a power law, the distribution of the number of Tweets shows deviations from the power-law distribution, while the  $R^2$ -value of fitting an exponential function  $y(x) = ae^{-bx}$  is significantly lower, supporting that a pure exponential fit is less appropriate. Functions of the type  $y(x) = \beta x^{-\alpha} + ae^{-bx}$ , which induce an exponential cut-off of the power-law's long tail, qualitatively show a better congruence with the distribution and intuitively correspond to the reasonable assumption that extremely large sizes of information spaces of users in SN and MSN platforms are very rare.

*Sub-Figure 3(c)* shows the distribution of spatial (geodesic) distance between adjacent nodes (actors with a direct social relation) in the social network. Equivalence classes of geodesic distances are determined in steps of 10 km. Due to the spherical topology of earth's surface (with a maximum circumference of roughly 40000 km at the Equator), the maximum class of spatial distances encompasses all geodesic distances between 19990 km and 20000km. As reasonably expected, the distribution shows two users with a smaller spatial distance have a higher probability of being socially connected, where the distribution roughly follows a power law. This confirms other study's results, such as [34] and supports the assumption that social closeness and spatial closeness mutually imply each other to a certain extent. As the diagram depicted in the left corner of the diagram shows, the geographic distribution of the users concentrates on the densely populated areas of North America and Europe. The dip of the curve around  $\approx 5000\text{km}$  may be explained by the relative geometric dimensions of the Atlantic ocean and the North American and European continent.

*Sub-Figure 3(d)* shows the correlation between the spatial distance of pairs of users (this time counted in classes of 50 km steps) and the semantic similarity of their information spaces (counted in equivalence classes of 1 %). The semantic



**Figure 3: General properties of the dataset and mutual implication of social, semantic and spatial closeness using different measures (compare discussion in the text) [14]. Wherever a curve fit is provided (e.g. a power law, linear or logarithmic function), standard regression [45] is used where  $R^2 = 1 - \sum_i (y_i - f(x_i, \beta))^2 / \sum_i (y_i - \bar{y})^2$  is the coefficient of determination [14].**

similarity of information spaces was computed as the Tanimoto coefficient [36] of the multi-set of term-frequency vectors of the respective sets of information items. Other alternatives would have e.g. been to use Rocchio centroids [27]. As an implementation, we used Lucene [1]. Of a matrix containing the absolute frequency of occurrences for a combination of a geodesic distance class and class of semantic similarity of information spaces we computed the average absolute frequencies for the four new equivalence classes [0, 100km], [100, 1000km], [1000, 5000km], and [5000, 20000km]. The four qualitatively Gaussian curves show that for larger distances the semantic similarity of the information spaces of the users is smaller than for smaller distances. This supports the connection between semantic relatedness and geographic relatedness. Qualitatively similar results have been obtained by [22] although the measures used were different.

*Sub-Figure 3(e)* depicts a correlation between the semantic similarity of information spaces of users (computed as in sub-figure 3(d)) and their average path distance in the social network. (Technically: of a matrix containing the absolute frequency of occurrences for a combination of a class of semantic similarities between  $[x, x + 1]\%$  and a path distance in the social network, we computed for each class of semantic similarities between  $[x, x + 1]\%$  the average over all path distances between 0 and 25). The result shows that the more similar the information spaces the smaller is the average social distance between the respective users. This supports the correlation between social closeness and semantic closeness.

*Sub-Figure 3(f)* shows a correlation between the geographic similarity of information spaces of users and their semantic similarity. While semantic similarity was computed in the same way as in 3(e) and 3(d), the geographic similarity of information spaces of users was computed in the following way: In order to compute a spatial relevance density for the information space of a user, a point-like spatial reference  $\mu = (\mu_1, \mu_2)$  of an information item was transformed into a Gaussian density contribution  $\mathcal{N}(\mu, \sigma)(x)$  with diagonal sigma corresponding to a 500 km circle, cut off at  $|x - \mu| = 500$  km with the help of ArcGis [2]. All contributions (which properly respected the spherical geometry of earth's surface) were added to yield a user  $u_i$ 's spatial relevance density  $\rho_i(x)$ . The geographic similarity  $\text{sim}(u_i, u_j)$  of the information spaces of two users  $u_i$  and  $u_j$  was computed via a Jaccard-like measure:

$$\text{sim}(u_i, u_j) = \int d^2x \frac{\min(\rho_i(x), \rho_j(x))}{\max(\int d^2x \rho_i(x), \int d^2x \rho_j(x))} \quad (1)$$

Although most of the information spaces had a similarity of 0 (this large contribution was left out of the diagram) the values show a trend that the closer the geographic similarity of information spaces, the larger the semantic similarity. Although the slope of this trend is rather small, this finding supports the correlation between geographic reference of information spaces and their semantic similarity.

Relating this geographic similarity of information spaces to the spatial geodesic distance between users as shown in *sub-figure 3(f)*, yields a logarithmic trend supporting the reasonable connection that spatial closeness of users also implies similarity in the spatial references of their information spaces.

*Sub-Figure 3(h)* relates the social similarity between users computed as the Jaccard-index of the sets of friends of two users and the respective average semantic similarity of infor-

mation spaces (where the semantic similarity of information spaces is computed as in sub-figures 3(d), 3(e), and 3(f)). We see a power law relating the two quantities: the more socially similar two users are, the more similar are their friend-sets and vice versa. This supports the connection between social and semantic contexts.

These preliminary results are an excellent ground for future research, investigating the connections between social, spatio-temporal and semantic contexts.

#### 4.4 Information Retrieval Experiments

The results just discussed show that the dataset can be viewed as a dataset *realistically* including and relating social, spatial and semantic elements. They *support the basic findings* of subsection 2.1 and subsection 2.2 and the *grounds* for the IR approach discussed in section 3. In order to *evaluate the basic suitability of these connections for IR*, IR experiments were conducted with the data-set.

As *queries*, Tweets were used. In the absence of real user assessments of relevance to be used as *ground truth* for the experiments, two *implicit assessments of relevance* were used as ground truths: As a *first* assessment of relevance, the *Re-Tweets* of the query Tweet were regarded as relevant. This assessment of relevance is intended to represent relevance with respect of the *conscious information need* of users. As a *second* assessment of relevance, all Tweets and Re-Tweets of users *following* (see [4]) the author of the query Tweet were regarded as relevant. This assessment of relevance is intended to represent relevance with respect to the *unconscious information needs* of users containing the contextual seeds discussed in subsection 2.1 and subsection 2.2.

In order to compare semantic search, social search and spatial search (excluding temporal aspects for reasons of simplicity) as a contextual bracket implicitly relating social and semantic contexts, *seven types of retrieval processes* were tested on the data-set. For each type of retrieval, the 50 best results (according to the IR model of the respective type) are retrieved and analyzed with the *first (I)* and *second (II)* 'ground truth' assessment of relevance by computing the usual confusion matrix (TP, FP, TN, and FN) and from that precision P and recall R [37]. If less than 50 items could be retrieved, either the missing ones are padded with random items from the respective pre-filtering (e.g. geographic or social) (*variant A*) before computing the measures to ensure comparability, or the measures are computed as is (*variant B*).

*Type 1 [Sem]:* semantic search (standard IR): Use Lucene [1] to compute a global IR index (over all information items of the dataset) and decide upon the 50 best matches to the query Tweet using Lucene's ranking.

*Type 2 [Soc]:* social search (social pre-filtering and subsequent semantic filtering): Retrieve all information items authored by friends and friends of friends of the query Tweet's author, compute a local IR index on these items and decide upon the 50 best matches to the query Tweet using the local index. This type of search is roughly associated with the expert-link-based type of social search with subsequent evaluation using a local IR system in the architecture.

*Type 3 [Geo]:* geographic search (geographic pre-filtering and subsequent semantic filtering): Using our implementation of our variant of distributed Quad-Tree and an octagonal query geometry centered around the query Tweet's spatial point reference of 'radius' between 500 km and 20

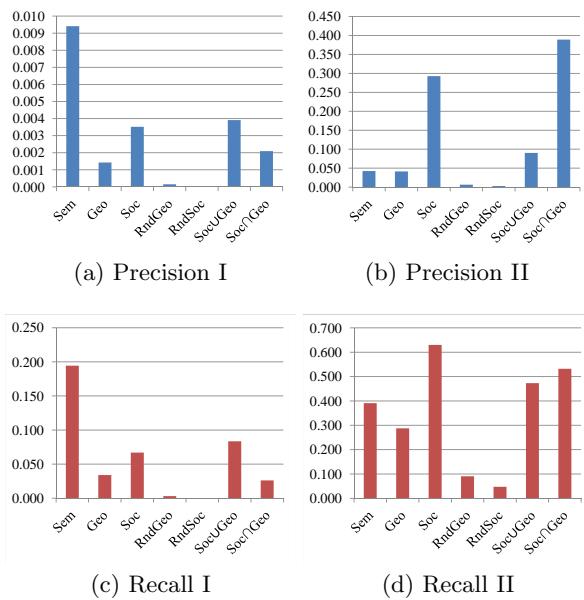
km depending on the depth of the tree in this region (corresponding to the density of information items), the spatially matching items were retrieved. On this set of items the semantically 50 best were determined as in the case of social search. This type of search is roughly associated with the spatio-temporal search of the architecture on Expertises with a subsequent employment of local IR.

*Type 4* [ $\text{Soc} \cup \text{Geo}$ ]: social-geographic search  $\cup$  (using the union  $X \cup Y$  of the results of geographic  $X$  and social pre-filtering  $Y$  and subsequent semantic filtering with Lucene as in type 2 and 3). This type is roughly associated with the spatio-temporal search of the architecture on all knowledge flags (Expertises and Expert-Links) with subsequent local IR.

*Type 5* [ $\text{Soc} \cap \text{Geo}$ ]: social-geographic search  $\cap$  (using the intersection  $X \cap Y$  of the results of geographic  $X$  and social pre-filtering  $Y$  and subsequent semantic filtering). This type of search is performed for reference purposes.

*Type 6* [ $\text{RndGeo}$ ]: random pre-filtering geographic (randomly select as many items from the dataset as a geographic pre-filtering would deliver and perform subsequent semantic filtering). This type of search is performed for reference purposes to further investigate the impact of geographic pre-filtering and thus the role of spatial context as a contextual bracket.

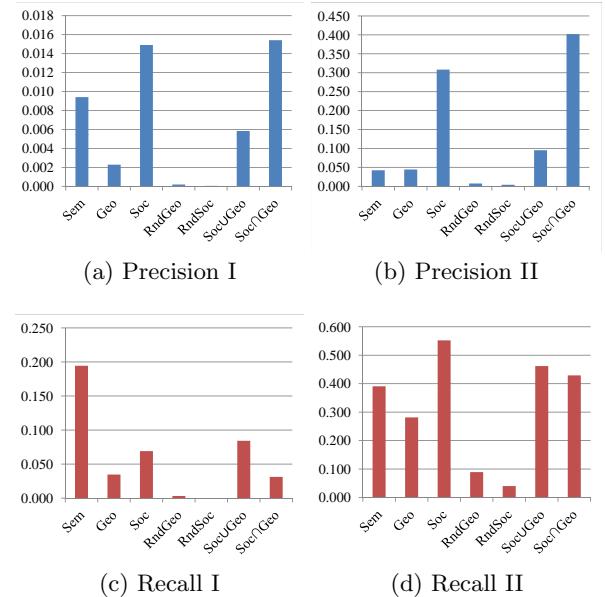
*Type 7* [ $\text{RndSoc}$ ]: random pre-filtering social (randomly select as many items from the dataset as a social pre-filtering would deliver and perform subsequent semantic filtering). This type of search is performed for reference purposes to further investigate the impact of social pre-filtering.



**Figure 4: Precision and Recall, variant A [14]**

Figure 4 shows the precision and recall values of variant A. The while for I, the first way of ground truth relevance assessment, the conventional purely semantic search performs best by far (in precision as well as recall), social search is most successful for II, the second way of ground truth relevance assessment and geographic search is still comparable to semantic search. If the assumption that II corresponds to

contributing to satisfying unconscious information needs via contextual seeds is indeed substantial, this result supports the proposed IR approach. In view of the role of spatial context as a context bracket implying semantic context to a certain degree, the comparison of the performance of geographic search (Geo) compared to random pre-filtering geographic (RndGeo) shows that indeed, Geo is significantly better than RndGeo. In other words, while Sem may use the whole set of information items to choose the 50 best (via the global index), Geo must choose from the considerably smaller set resulting from geographic pre-filtering and still delivers acceptable relative performance compared to a random pre-filtering. Indexing the whole set of information items may not be desirable for SN and MSN environments due to privacy considerations. Because of the connections between geographic closeness and social closeness, we can thus, in a realistic SN and MSN setting, expect that Geo may effectively draw from a locally richer set of relevant items and thus deliver even better overall performance than Sem.



**Figure 5: Precision and Recall, variant B [14]**

Figure 4 shows the precision and recall values of variant B, where a, due to the restrictive pre-filtering, insufficient number of retrieved items is not padded by random items (which induces a pessimistic evaluation for the contextual search variants). Here, as a consequence, social search is best also for assessment I with respect to precision.

## 5. CONCLUSION AND OPPORTUNITIES FOR FUTURE RESEARCH

Our *overall results* may be interpreted as giving *support* to exploiting the concept of Spatio-Temporal Small Worlds and the underlying correlations between semantic, spatio-temporal, and social contexts for alternative IR, akin to Human IR in (Decentralized) Social Networking.

However, the evaluation environment may still not take

advantage of several of the benefits of the architecture (such as the power of local agent IR systems). Thus, one might expect that the approach is indeed able to deliver useful contextual seeds especially in view of unconscious information needs and thus is a new alternative IR concept for SN and MSN environments.

Nevertheless, the introduced study is only a starting point for a large body of *future work* on connecting social, semantic and spatio-temporal contexts for new and useful forms of IR.

As has been mentioned above, a full implementation and real world evaluation of the architecture would be the next step following the usual Design Science methodology [25]. A special focus has to be put on evaluating the usefulness of the results obtained by the suggested alternative IR methods in terms of the extended notions of information need discussed above. Suitable concepts of extended versions of precision and recall will have to be constructed for the respective evaluations. Furthermore, more variants of combining spatial, social and semantic retrieval criteria need to be evaluated in relation to the individual and social short term and long term context of the querying user.

## 6. REFERENCES

- [1] Apache lucene search engine library, 2011. <http://lucene.apache.org/java/docs/index.html>, (checked May 2012).
- [2] Esri arcgis software system, 2011. <http://www.esri.com/software/arcgis>, (checked May 2012).
- [3] Metacarta geo-coding web-service, 2011. <http://www.metacarta.com/>, (checked May 2012).
- [4] Twitter platform. <http://twitter.com>, (checked May 2012), 2012.
- [5] R. Angelova, M. Lipczak, E. Milius, and P. Pralat. Investigating the properties of a social bookmarking and tagging network. *International Journal of Data Warehousing and Mining (IJDWM)*, 5(0):12–29, 2009.
- [6] N.J. Belkin. Interaction with texts: Information retrieval as information-seeking behavior. *Information Retrieval*, 93:55–66, 1993.
- [7] N.J. Belkin, R.N. Oddy, and H.M. Brooks. Ask for information retrieval: Part i. background and theory. *Journal of Documentation*, 38(2):61–71, 1982.
- [8] N.J. Belkin, R.N. Oddy, and H.M. Brooks. Ask for information retrieval: Part ii. results of a design study. *Journal of Documentation*, 38(3):145–164, 1982.
- [9] William B. Cavnar and John M. Trenkle. N-gram-based text categorization. In *In Proceedings of SDAIR-94, 3rd Annual Symposium on Document Analysis and Information Retrieval*, pages 161–175, 1994.
- [10] D. Centola. The spread of behavior in an online social network experiment. *science*, 329(5996):1194–1197, 2010.
- [11] D. Centola and M. Macy. Complex contagions and the weakness of long ties1. *American Journal of Sociology*, 113(3):702–734, 2007.
- [12] M.S. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380, 1973.
- [13] G. Groh, F. Straub, and B. Koster. Spatio-temporal small worlds for decentralized information retrieval in social networking. In *Proc. ACM SigSpatial 2012*. ACM, 2012.
- [14] Georg Groh. *Contextual Social Networking*. Habilitation Thesis, TU-München, 2012.
- [15] Georg Groh and Stefan Birnhammerer. Privacy and information markets: Controlling information flows in decentralized social networking. *Proc. IEEE Socialcom'11, Boston, USA*, 2011.
- [16] Georg Groh and Christian Ehmig. Recommendations in Taste Related Domains: Collaborative Filtering vs. Social Filtering. *Proc. Group07, Sunibel Island, USA, Nov 2007*, 2007.
- [17] Georg Groh, Alexander Lehmann, Tianyu Wang, Stefan Huber, and Felix Hammerl. Applications for social situation models. *Proc. Int'l Conf. Wireless Applications and Computing Conference, Freiburg, Germany 2010*, 2010.
- [18] Georg Groh and Florian Straub. An architecture for an alternative, multi-agent-based information retrieval approach with spatio-temporal primary classification criterion. *GIS.Science Journal*, 02/2010, 2010.
- [19] Georg Groh, Florian Straub, Andreas Donaubauer, and Benjamin Koster. Space and time as a primary classification criterion for information retrieval in distributed social networking. *Arxiv publication, online via http://arxiv.org/abs/1104.2196*, 2011.
- [20] B. He, M. Patel, Z. Zhang, and K.C.C. Chang. Accessing the deep web: a survey. *Communications of the ACM*, 50(5):95–101, 2007.
- [21] B. Hecht and E. Moxley. Terabytes of tobler: Evaluating the first law in a massive, domain-neutral representation of world knowledge. In *Proceedings of the International Conference on Spatial Information Theory (COSIT 2009)*, pages 88–105. Springer, 2009.
- [22] B. Hecht and M. Raubal. Geosr: Geographically explore semantic relations in world knowledge. *Proc. 11th AGILE International Conference on Geographic Information Science, Girona, Spain (2008)*, pages 95–114, 2008.
- [23] B. Hecht and J. Schöning. Mapping the zeitgeist. In *Proceedings of the 4th International Conference on GIScience, Extended Abstracts*, 2008.
- [24] B.J. Hecht and D. Gergle. On the localness of user-generated content. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, pages 229–232. ACM, 2010.
- [25] A.R. Hevner, S.T. March, J. Park, and S. Ram. Design science in information systems research. *Mis Quarterly*, pages 75–105, 2004.
- [26] S. Jin and A. Bestavros. Small-world characteristics of internet topologies and implications on multicast scaling. *Computer Networks*, 50(5):648–666, 2006.
- [27] T. Joachims. A probabilistic analysis of the rocchio algorithm with tfidf for text categorization. In *ICML '97 Proceedings of the Fourteenth International Conference on Machine Learning*, pages 143–151, 1997.
- [28] C. Jones and R. Purves. Geographical information retrieval. *International Journal of Geographical Information Science*, 22(3):219–228, 2008.
- [29] J. Kleinberg. The small-world phenomenon: an

algorithmic perspective. In *Annual ACM Symposium on Theory of Computing*, volume 32, pages 163–170, 2000.

[30] J.M. Kleinberg. Navigation in a small world. *Nature*, 406(6798):845–845, 2000.

[31] Benjamin Koster. *Simulation and Evaluation of an Approach for Federated, Agent-based, Spatio-temporal Information Retrieval*. Bachelor’s Thesis, TU-München, WS 2009 / 2010; Supervisor: Georg Groh (co-supervised with Florian Straub), 2010.

[32] R.R. Larson. Geographic information retrieval and spatial browsing. In *Smith, Linda C., Gluck, Myke (eds.): Geographic information systems and libraries: patrons, maps, and spatial information : Papers presented at the 1995 Clinic on Library Applications of Data Processing, April 10-12, 1995*, pages 81–124. University of Illinois, 1996.

[33] D. Liben-Nowell. Wayfinding in social networks. In *G. Cormode, M. Thottan (eds.): Algorithms for Next Generation Networks*, pages 435–456, 2010.

[34] David Liben-Nowell, Jasmine Novak, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Geographic routing in social networks. *Proceedings of the National Academy of Sciences*, 102(33):11623–11628, August 2005.

[35] M.D. Lieberman and J. Lin. You are where you edit: Locating wikipedia contributors through edit histories. *Proceedings of ICWSM09*, 2009.

[36] A.H. Lipkus. A proof of the triangle inequality for the tanimoto distance. *Journal of Mathematical Chemistry*, 26(1):263–265, 1999.

[37] C.D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.

[38] M. McPherson, L. Smith-Lovin, and J.M. Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, pages 415–444, 2001.

[39] S. Milgram. The small world problem. *Psychology Today*, 2(1):60–67, 1967.

[40] S. Mizzaro. How many relevances in information retrieval? *Interacting with Computers*, 10(3):303–320, 1998.

[41] D. Nadeau and S. Sekine. A survey of named entity recognition and classification. *Lingvisticae Investigationes*, 30(1):3–26, 2007.

[42] M.E.J. Newman. The structure and function of complex networks. *SIAM review*, pages 167–256, 2003.

[43] J. O’Madadhain, D. Fisher, S. White, and Y. Boey. The jung (java universal network/graph) framework, 2003. <http://jung.sourceforge.net>, (checked May 2012).

[44] M.F. Porter. An algorithm for suffix stripping. *Program: Electronic Library and Information Systems*, 14(3):130–137, 1980.

[45] H. Pruscha. *Statistisches Methodenbuch: Verfahren, Fallstudien, Programmcodes*. Springer Verlag, 2005.

[46] S. Scellato, A. Noulas, R. Lambiotte, and C. Mascolo. Socio-spatial properties of online location-based social networks. *Proceedings of ICWSM*, 11:329–336, 2011.

[47] X. Shen, B. Tan, and C.X. Zhai. Context-sensitive information retrieval using implicit feedback. In *Proc. 28th Int’l. ACM SIGIR Conf. on Research and Development in Information Retrieval*, pages 43–50. ACM, 2005.

[48] C. Tang and S. Dwarkadas. Hybrid global-local indexing for efficient peer-to-peer information retrieval. In *Proc. 1st Symposium on Networked Systems Design and Implementation*, pages 16–16. USENIX Association, 2004.

[49] W.R. Tobler. A computer movie simulating urban growth in the detroit region. *Economic Geography*, 46:234–240, 1970.

[50] D.J. Watts and S.H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, 1998.

[51] C. Wilson, B. Boe, A. Sala, K.P.N. Puttaswamy, and B.Y. Zhao. User interactions in social networks and their implications. In *Proceedings of the 4th ACM European Conference on Computer Systems*, pages 205–218. ACM, 2009.

[52] T.D. Wilson. Human information behavior. *Informing Science*, 3(2):49–56, 2000.

[53] C.A. Yeung, I. Liccardi, K. Lu, O. Seneviratne, and T. Berners-Lee. Decentralization: the future of online social networking. In *W3C Workshop on the Future of Social Networking Position Papers*, 2009.